

MASSTOR SYSTEMS

M860 STORAGE MANAGEMENT SYSTEM

System Description

Strategic Solutions

Masstor Systems' strategic product set provides integrated systems solutions for Moving, Storing and Sharing large amounts of data. New strategic products and solutions are required to manage the continued proliferation and dispersment of data processing power with the need to provide rapid access to large amounts of data.

- In the 1960's, the computer power was concentrated in the "machine rooms" and access to data was easy.
- In the 1970's, the computer power became more dispersed. The power outside the "machine room" equalled the power within the "machine room" and access to data became more complex.
- By the 1990's, the computer power outside the "machine room" will have grown to at least 10 times that within the "machine room." The "machine room" now has many mainframe computers, and access to data is becoming very complex.

Masstor Systems is committed to the enhancement of its product set to provide solutions for the environment of the 1990's. MASSNET for Moving, M860 Storage Management System (SMS) for Storing, and Shared VSS for Sharing access to large amounts of data.

MASSNET Moves Data

MASSNET is a family of hardware and software networking products that moves data at high speeds, up to 1000 times faster than telecommunication-based networks. MASSNET is used to move data among local or remote mainframes of similar and dissimilar manufacture.

M860 SMS Stores Data

M860 Storage Management System is a very high capacity on-line storage device that attaches to IBM or IBM Plug Compatible Manufacturer's mainframes. The M860 SMS improves the economies of storing large amounts of data on-line, and it stores up to 18 times more data than conventional disk in an equivalent amount of floor space.

SHARED VSS Shares Data

Shared Virtual Storage System is an integrated system solution for controlling and managing large amounts of data being shared in a dispersed computing environment. Shared VSS incorporates MASSNET and the M860 SMS and extends the benefits of the M860 to multiple mainframes of similar and dissimilar manufacture.

Summary

Our product solutions target the emerging needs of the world's largest data processing organizations. A large number of major corporations and institutions have installed Masstor Systems' products in broadcasting, banking, insurance, aerospace manufacturing, telephone, railroad, airline, oil, government, university and research environments.

Users of the Masstor Systems' strategic product set have realized the benefits of increased levels of productivity and service by implementing new applications such as automatic Disaster Backup and Recovery with economic benefit to the user. These applications dramatically increase the ability to share resources and improve the access to data.

MASSTOR SYSTEMS

Strategic Solutions to the Problems
of Information Management

Table of Contents

	PAGE
INTRODUCTION	1
The Demand For Mass Storage	1
DISK STORAGE	1
Disk—A Jack of All Trades	2
Specialized Companions to Disk	2
Disk Space Management Problems	2
Disk Space Management Systems	2
THE M860 STORAGE MANAGEMENT SYSTEM (SMS)	2
Assured Data Integrity	2
Automated Operation	2
USES OF AUTOMATED, LARGE CAPACITY MASS STORAGE	2
Tape Replacement	2
Disk Space Management	2
COMPONENTS OF THE M860 SMS	4
The M860 Storage Management Task (SMT)	4
The M861 Storage Module	4
Model and Feature Summary	4
The Data Recording Complex	4
Data Cartridge	4
Data Recording Device	4
Cartridge Storage Complex	4
Storage Cells	5
Manual Cartridge Access Station	5
Accessor	5
THE M862 STORAGE CONTROLLER	5
System and Device Attachment	5
Function	5
Features	6
USES OF THE M860 SMS	6
A Disk Alternative	6
A Tape Library Alternative	6
A Combination	7

Introduction

The Masstor Systems M860 Storage Management System (SMS) is a compact high-capacity direct-access magnetic data store. The M860 SMS addresses modern data processing requirements for vast storage capacity, guaranteed data integrity, unattended operation and rapid access to data. This brochure reviews the current data processing trends Masstor Systems addresses with the M860 SMS; it provides an overview of the

components and features of this unique storage device; and it describes the role the M860 SMS plays in solving data storage problems.

The Demand for Mass Storage

The three-decade emergence of data processing can be depicted as a spiraling growth of computing power and demand for data storage capacity. Every improvement in computer hardware spawns new developments in operating

systems and application bases. Every application improvement generates demand from the user community for additional refinements, faster computers and GREATER STORAGE CAPACITY.

Disk Storage

For the mainstream data processing installation with typical applications, progress in disk storage technology has satisfied the majority of the on-line storage requirements. The familiar storage media to

the systems analyst and to the applications programmer is disk storage—because every data processing configuration includes that quintessential storage device.

Disk is an acceptable storage medium because it satisfies a wide variety of storage requirements. Disk reliability, density, speed of access and cost per byte today permit many applications to utilize on-line storage capabilities. A software base has developed around these



Figure 1: MASSTOR M860 Storage Management System

familiar storage devices, facilitating implementation of new applications using disk. Most of the indices of disk cost-effectiveness keep improving, but there are problems with disk that have traditionally compelled users to supplement disk storage with specialized mass storage devices.

Disk—A Jack of All Trades

A disk drive has the flexibility to hold a vintage ISAM file, a program library, a VSAM database, an IMS log and a system paging dataset. Barely fast enough to fulfill the paging function, the disk is not cheap enough to use as a mass data store. For all these applications disk reliability is acceptable, but not good enough to use disk as the final repository for any item of data. It is a fact of life, hard-learned at less fortunate installations, that all modern disk devices incur an exposure to catastrophic and permanent media destruction.

Specialized Companions To Disk

Earlier mass storage devices were designed to address a particular level of performance in terms of per-byte cost and data transfer capacity. The segment of the marketplace that required those levels of performance accepted and used those devices. Today the industry's demand for storage density and data delivery rate has gone beyond what those devices were designed to provide, just as it has surpassed the capabilities of the earlier disk drives such as the IBM 2314, the IBM 3330 and the IBM 3350.

Today a single device, the IBM 3380 disk subsystem is required to fulfill every disk storage and mass storage requirement generated by a conventional IBM installation from MVS paging to archive storage. A similar problem exists in non-IBM installations, in many cases aggravated by even higher costs per-byte.

Disk Space Management Problems

Near-total reliance on a volatile medium such as disk raises the requirement for management of that disk resource to an urgent level of priority. This generates the need for disk backups, for performance tuning of the disk subsystem and for tracking of disk contents. The Data Processing center that uses on-line storage to provide uninterrupted service, fast response and unattended operation for its own customers must use third-generation utilities and manual tape backup procedures to tend the disk farm those customer applications create.

Disk Space Management Systems

Current versions of disk space management systems such as IBM's DFHSM, Sterling Software's DMS/OS or Cambridge Systems' ASM2 provide software tools to solve the disk problem. But history shows us the only way to alleviate the shortcomings of disk as a permanent storage device is to supplement disk with a specialized mass storage subsystem.

The M860 Storage Management System (SMS)

The Masstor Systems' M860 SMS consists of an M862 Storage Control, one or more M861 Storage Modules, and a Storage Management Task (SMT) software component. The M860 attaches to IBM-compatible block multiplexor channels using interlock or data-streaming protocol. The SMT is supported on any release of MVS/370 since MVS 3.7, including MVS/XA, and in any JES configuration.

The M860 SMS appears to the host operating system as a cluster of magnetic tape drives. Each M861 Storage Module contains the equivalent of two drives, which can emulate the characteristics either of 3420 series magnetic tapes or (as an option) 3480 cartridge tape drives.

Assured Data Integrity

Although it emulates existing IBM tape devices, the M860 implements a read-after-write discipline for every block written and fully automated error recovery. These integrity assurances, plus the elimination of physical tape handling, make the M860 a much more reliable data store than the tape devices it emulates.

Automated Operation

The Storage Management Task software coordinates allocation requests for logical tape volumes residing in the M860, and translates MVS mount messages into sequences of commands to the M860. Allocation and mounting of M860 volumes is accomplished transparently and automatically. The storage capacity of the M860 can be shared by up to four independent computer systems, under control of SMT. SMT provides fully automated 4-way sharing of M860 units, an extension of MVS function.

Uses of Automated, Large-Capacity Mass Storage Tape Replacement

The storage capabilities of the M860 SMS are readily available to any application that currently uses 3420 or 3480 tape.

Random-access applications that could never be implemented in a cost-effective manner on disk, run on M860s today.

The M860's tape-handling capabilities alleviate operator tape-handling responsibilities in a variety of applications for which response time, physical integrity and data security are of critical importance.

The M860 provides an economical, across-the-board replacement for a site's baseline tape or 3480 usage (i.e., the modal number of drives in use). The M860 is not a cost-effective replacement for seldom-used tape drives, because an idle M861 storage module is more expensive than two idle tape drives, and unused tape drives do not generate operator overhead.

Disk Space Management

The most popular use of the M860 SMS, a use for which the device is particularly well-suited, is as a migration ("archive") device for one of the Disk Space Management Systems (DSMS) described previously. The DSMS makes the M860 cost-effective by using it to off-load data from disk. This off-loading is achieved within the fixed overhead that must be allotted for managing the volatile disk media, through a technique called "incremental backup".

The combination of the DSMS and the M860 SMS trades an operator-intensive, uni-directional dump of disk to tape for an unattended off-loading of disk to cheaper safer M860 storage. Data sets off-loaded to M860 are automatically restored at rates of hundreds per hour. Previously, each data set recalled from the tape or 3480 library involved the efforts of a tape operator and minutes of elapsed time.

Disk space management techniques translate into optimal use of valuable disk resources as high-performance storage. They insure data against accidental permanent destruction. In the long term, a disk space management program with the M860 enables a data center to stop paying disk storage rates for data that does not require the performance characteristics provided by disk. The M860 SMS permits a site to implement a program of effective disk space management, without sacrificing rapid, unattended access to critical data.

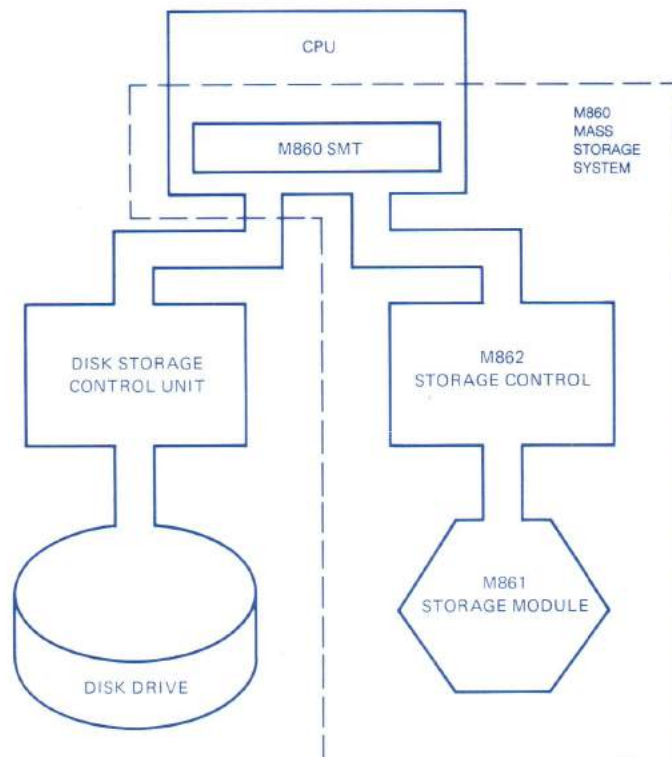


Figure 2: Components of the MASSTOR M860 Storage Management System




Figure 3: Major Components of the M861 Storage Module

Components of the M860 SMS

The major components of the M860 SMS are:

- The M860 Storage Management Task (SMT) host software.
- The M861 Storage Module.
- The M862 Storage Controller.

Figure 2 shows the components of a minimum configuration of the M860.

Up to four M861 Storage Modules can be attached to an M862 Storage Controller. Two M862 Storage Controllers and up to eight M861 Storage Modules can be coupled to provide redundant access paths to data. Multiple M862s and M861s can be attached to a host machine.

The M860 Storage

The M860 SMT host program is a task which executes under the IBM operating system (as described previously). It communicates with the M862 Storage Controller over decided control paths. Functions of SMT are:

- To maintain tables that show the location of every logical volume and every data cartridge within the M860.
- To communicate with other host processors in a multiple CPU environment.
- To send high level, logical requests to the M862 Storage Controller.
- To Service Requests from the M862 Storage Controller.
- To perform volume creation and volume deletion.
- To perform volume mount and volume dismount.
- To provide the interface to the SMS for operator commands and SVC requests of authorized users.
- To perform device allocation and de-allocation.
- To provide resource conflict resolution.
- To perform certain error recovery and alternate path retry operations.

The M861 Storage Module

The M861 Storage Module consists of two functional components that work

together to provide the facilities of the M860 Storage Management System. They are (1) the Data Recording Complex, and (2) the Cartridge Store Complex. Figure 3 shows most of the major components of the M861 Storage Module.

Data is stored on tape cartridges that reside in 316 cartridge cells. When data is written or read, the accessor moves to the appropriate cell, removes the cartridge from the cell, and transports it to a Data Recording Device (DRD). The cartridge is then loaded into the DRD, which writes or reads data to or from the tape. The data is transmitted through the M862 Storage Controller to or from a channel on the host Central Processing Unit. After the requested data has been written or read, the data cartridge is returned to its home cartridge cell.

Model and Feature Summary

There is one model of the M861 Storage Module. It contains the Data Recording Complex and the Cartridge Store Complex. The Data Recording Complex consists of data cartridges and two Data Recording Devices. The Cartridge Store Complex is composed of cartridge storage cells, a cartridge access station, and an accessor.

The Data Recording Complex

The two components of the data recording complex perform the transfer of data to and from the Central Processing Unit.

Data Cartridges

The M860 Storage Management System uses data cartridges on which to store data. Each data cartridge contains a length of magnetic tape that is wound on a plastic spool and enclosed in a plastic, protective shell.

Data is contained on the magnetic tape in the image of tape volumes. Multiple volumes of variable length can be recorded onto the same data cartridge.

Data is recorded on the tape serially, bit by bit, in a series of diagonal paths, called stripes, which are numbered beginning with zero. Each data cartridge contains 27,088 stripes. The first 40 stripes hold the Cartridge Table Of Contents (CTOC). The remaining 27,048 stripes are available for user data.

Each stripe can contain up to 6652 bytes of data. The first 32 bytes and the last 32 bytes are reserved for internal use of the system for error checking and recovery. Each user data block or portion of a block within a stripe contains a 16-byte header and a 16-byte trailer. In addition, a Cyclic Redundancy Check (CRC) byte is inserted after every 254 user data bytes. Therefore, the maximum number of user data bytes per stripe is 6530.

The maximum user data capacity per data cartridge is 176,623,440 bytes (27,048 stripes x 6530 bytes/stripe).

Each stripe carries its own identification, which is called its "stripe number". Servo information along the outside edges of the tape allows the data recording device to position to the correct stripe.

Data Recording Device

A Data Recording Device reads and writes the data stored in the M860 SMS's data cartridges. It has three cartridge stations: (1) an entry station, to which the accessor delivers the cartridge to be read or written, (2) a read/write station, where the cartridge remains while the Data Recording Device is either reading from it or writing onto it, and (3) an exit station, from which the accessor retrieves the cartridge to return it either to its home cell or to deliver it to the exit cell of the cartridge access station.

When the Data Recording Device is empty and a data cartridge has entered the entry station, the cartridge is moved to the read/write station, and the spool of tape is removed from its protective outer shell. The tape is threaded through a read/write transport, around a two-part mandrel (a stationary

cylinder) in a helical path, and onto a take-up spool. The helical path taken by the tape and the location of the read/write heads between the mandrel halves are shown in Figure 4. The figure also shows how the stripes are written diagonally on the tape.

Read and write heads on a rotor in the mandrel write data and perform a read-back check of the written data on the same revolution of the rotor. If the user volume requested is not at or near the beginning of the tape, the Data Recording Device performs a high-speed search to the required volume. When the M860 is defined to the host system as an IBM 3480 tape system, the M860 performs a high-speed search at 36 times the read/write speed to the block of data required. Once the tape is positioned, it is held stationary during the reading or writing of each stripe. When the heads reach the end of a stripe, the tape is moved to the next stripe. Whenever the tape is physically in motion, it rides on an air bearing to eliminate wear.

When all the activity scheduled for the currently loaded cartridge is complete, the M862 Storage Controller initiates a Search order, which causes the Data Recording Device to search to the beginning of the tape. The M862 Storage Controller then updates the cartridge's label information and initiates an Unload order which causes the Data Recording Device to replace the tape spool in its protective shell and insert the cartridge into the exit station. The accessor then removes the cartridge from the exit station and returns it to its storage cell.

During each unloading operation, the read/write heads are automatically cleaned by a head cleaner that is built into the Data Recording Device.

Cartridge Store Complex

The Cartridge Store Complex contains the storage cells, the manual cartridge access station, and the accessor. It accepts commands from the M862 Storage Control and controls the data cartridge




Figure 4: Mandrel of the Data Recording Device

movement within the M861 Storage Module based upon those commands.

Storage Cells

Data cartridges are stored in the hexagonal cells of a honeycomb-like wall which is shaped cylindrically around the two side walls and the back wall of the M861 Storage Module. Each storage cell has a unique physical address that indicates how high it is above the bottom row of cells (its Y coordinate), and how far it is from the left-most column (its Z coordinate).

Each M861 Storage Module has 320 cells. Four of them are reserved for use by the M860 Storage Management System. Therefore, 316 cells are available for storage of user data.

Manual Cartridge Access Station

Each M861 Storage Module has a manual cartridge access station which is used for inserting and removing data cartridges. It consists of an entry cell at the bottom of an entry chute, and below it, an exit cell at the top of an exit chute. The entry and exit cells each have YZ addresses in the

same format as those of the storage cells. Figure 5 shows the manual cartridge access station.

When a cartridge is placed in the entry chute, it rolls down to the entry cell location and trips a switch that signals the M860 SMT, through the M862 Storage Controller, that a cartridge is present. The M860 SMT then instructs the M862 Storage Controller to move the cartridge to a Data Recording Device. The Data Recording Device reads the cartridge label information and transmits it back through the M862 Storage Controller to the M860 SMT, which instructs the M862 Storage Controller to move the cartridge to a specific cell and saves the cartridge label information and the location of the cell for future reference.

For normal removal of cartridges from the M861 Storage Module, the M862 Storage Controller, on instructions from the M860 SMT, causes the accessor to move a cartridge to a Data Recording Device so that its label can be verified. Then the accessor moves the cartridge to the exit cell, where it rolls down the exit chute for manual removal.

Accessor

Each M860 Storage Module has an accessor and its controlling microcode (known as the accessor control) which moves data cartridges among the cell locations, the Data Recording Devices, and the cartridge access station.

A cartridge picker, which is located on the accessor, travels vertically in what is called the "Y" direction. The picker can also swivel in a circular motion around the side and back walls in a movement said to be in the "Z" direction. Figure 6 shows an M861 Storage Module accessor.

When a data cartridge must be moved, the accessor control causes the accessor to move from its current location to the location of the data cartridge. Then the cartridge picker retrieves the data cartridge, and the accessor delivers it to the location specified by the accessor control.

The M862 Storage Controller

The M862 Storage Controller provides the logical capabilities necessary to operate and control the attached M861 Storage Module(s).

System and Device Attachment

The M862 Storage Controller can attach to all IBM or other manufacturer's IBM Plug Compatible Central Processing Units which support the IBM OS/VS2 MVS Release 3.7 or higher operating systems including MVS/XA.

Up to four M861 Storage Modules can be attached to a single M862 Storage Controller.

Function

Major functions of the M862 are:

- To provide a data path and a control path between the Central Processing Unit and the M861 Storage Module.
- To communicate with the M860 SMT operating in the host computer.
- To execute the host generated Channel Commands.
- To perform the format conversion between user data blocks and stripes.
- To handle all operations of the accessor(s) and Data Recording Device(s).

Features

Each M862 Storage Controller is equipped with four channel attachments and two Data Recording Controllers (DRCs). Each DRC can operate and control up to eight Data Recording Devices (DRDs). Standard features include:

- **Four Channel Access**—Four Channel Access allows the M862 Storage Controller to be shared by four channels. The channels may be attached to either the same or different central processing units.
- **Dual Data Path**—The Dual Data Path feature provides a second path so that two independent transfers between the CPU(s) and/or M861 Storage Module(s) can occur simultaneously.

Uses of the M860 SMS

For users of the IBM or IBM/PCM CPUs, the directly attached M860 SMS provides

economical, on-line, large capacity data storage. Users of non-IBM type processors, such as Univac, Honeywell, or others can use the M860 SMS through the services of Masstor Systems' Shared Virtual Storage System (Shared VSS). Regardless of the CPU type or the method of attachment, the M860 Storage Management System can enhance an installation's operational and cost effectiveness as:

- A Disk Alternative
- A Tape Library Alternative
- A Combination

A Disk Alternative

In medium-to-large computer installations, studies have shown a measured growth of installed disk capacity in excess of 45 percent compounded annually. Most sources predict that this growth rate will not only continue, but more likely increase. Because improvements

in disk technology have not kept pace with the demand for on-line storage, many users are faced with rapidly increasing on-line storage costs and a severe shortage of computer room floor space.

A Tape Library Alternative

The M860 SMS appears to the operating system as though it were a tape subsystem (3420 or 3480), and supports all tape commands. As a result, the M860, for certain applications, can be used exactly as a tape replacement where one DRD is equivalent to one tape drive. An example of this type of application is the continuous high volume storage of geological or satellite data. Another example would be the backup and restoration of disk volumes. Again, no conversion or system changes are needed for this use.

The more typical use of the M860 will be as an alternative storage for the data currently stored on off-line reels of tape in a library. To maximize the productivity of the M860, the data set is first copied to disk where it is processed as a disk resident data set. This staging and destaging can be accomplished "explicitly" or "implicitly".

Study results from earlier versions of mass storage systems (MSS) have proven the dramatic benefits of this use. In many cases, if only 20 percent of the tape data is stored on MSS, tape drives and tape reels have been cut in half and tape mounts have been reduced by 70 percent. Aside from the direct cost savings realized, users also realize overall system throughput increases due to the elimination of errors and delays due to manual tape handling.




Figure 5: Manual Cartridge Access Station

A Combination

Most users will achieve the maximum benefits from the M860 Storage Management System through a combination of the two uses already described.

Today, however, the user has an alternative. An M860 Storage Management System can solve the problems of capacity, storage cost, and floor space.

Certain types of data sets, such as batch or infrequently used production data sets, convenience data sets, and development data sets are ideal candidates for storage on the M860 SMS. These data sets can account for over half of all the data currently stored on disk. When required for processing, the data sets can be automatically "staged" to

disk in a matter of seconds and processed at normal disk speeds.

Since these data sets are currently disk resident, no modifications to problem programs are necessary. The data sets are "archived" to the M860 and "recalled" from the M860 using the same procedures that would be used if the M860 was a bank of tape drives. It can be accomplished "explicitly" through standard IBM or other tape-to-disk (disk-to-tape) utilities. Or, the staging and destaging can occur "explicitly" or "implicitly", via an existing Disk Space Management System such as IBM's DF/HSM, Sterling Software's DMS/OS or Cambridge Systems Group's ASM2.




Figure 6: The Accessor